



Scientific Computing at Warwick:

Current facilities and developments

WARWICK


Dr David Quigley

Director, Centre and RTP for Scientific Computing

Facilities

Linux-based compute clusters

Cluster computing



- ▶ Multithread parallelism within a node (e.g. OpenMP)
- ▶ Message passing parallelism between nodes (MPI)
- ▶ Different workloads
 - Many nodes working together with frequent messaging over fast network
 - Many nodes working independently, i.e. task farming

Minerva

- ▶ Regional High Performance Computer (HPC)
- ▶ Installed in April 2012
- ▶ End of life

15/16 usage





■ Warwick ■ QMUL ■ Nottingham ■ Birmingham



- Around 2500 cores
- 89 Warwick users in 15/16
- Mainly phys, chem & eng
- 86 TB of data in /home

Apocrita


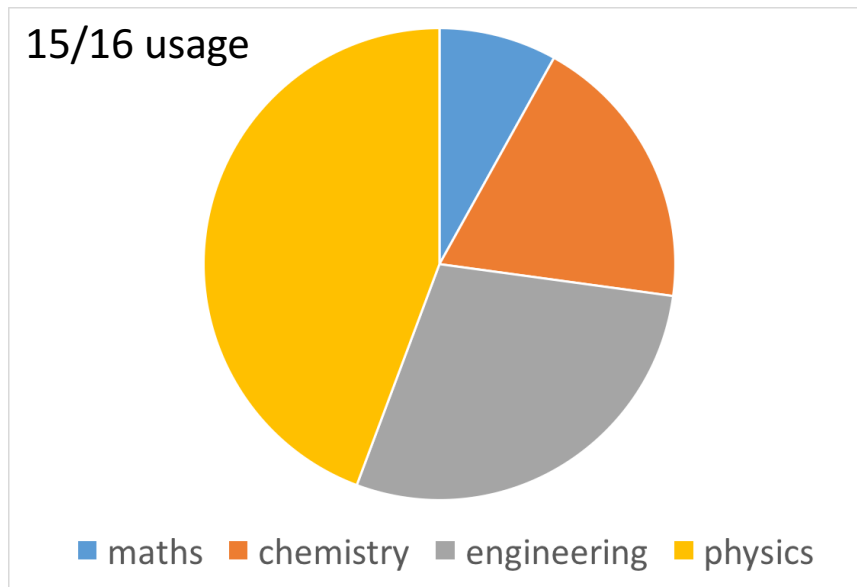
- ▶ Regional High Throughput Computer (HTC)
- ▶ Installed 2012
- ▶ End of life



- 2328 cores
- Warwick usage ~ 15%

Tinis

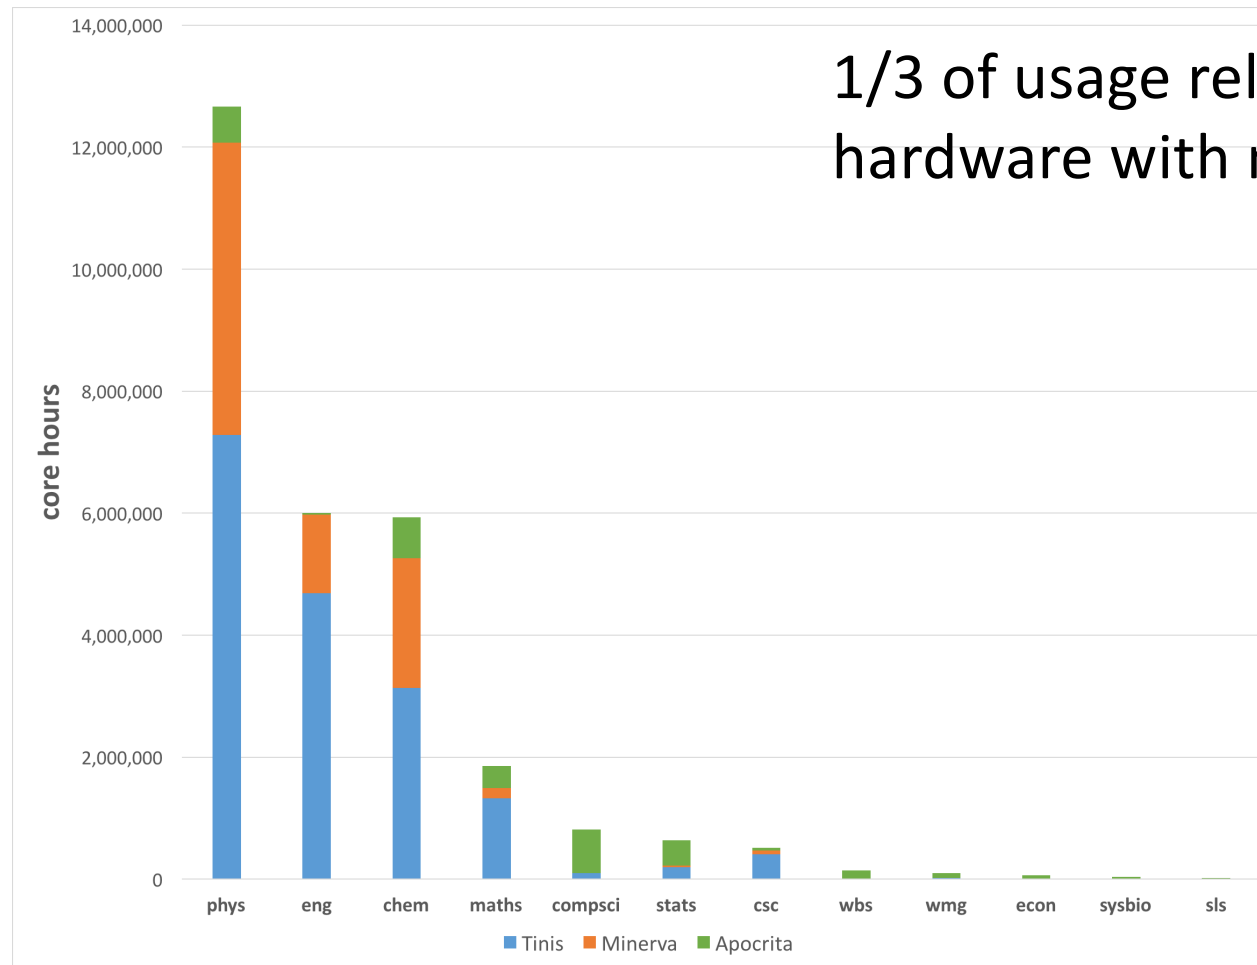
- ▶ Available since November 2015
- ▶ Not 100% new kit
- ▶ Expected 4 year life



- Purchase cost of £0.9M
- 3428 cores in 16-way nodes
- 1 PB of shared storage
- Some GPU & Xeon Phi nodes



Overall usage



1/3 of usage relies on end-of-life hardware with no service contract!



New cluster (2016/17)

- ▶ Requirements discussed with user board over last summer
- ▶ Will go out to tender shortly (£0.75M)
- ▶ Expecting to be similar to Tinis, but with latest generation of interconnect
- ▶ Actual capacity not yet known
- ▶ **There will be a substantial gap between decommissioning of Minerva and availability of the new system**



Usage reports

- ▶ Requested via the user board
- ▶ Not a bill or an invoice
- ▶ There are for information only

```
# Tinis usage report for Aug 2016.
# User      Share      Supervisor      Time Used (CPU hours)      Share Used (%)      Realname
phsdtt      physicsdq      David Quigley      58879.46              52                  Andrii Vasylenko
phrich      physicsdq      David Quigley      42629.62              38                  Samuel Brown
phrlbm      physicsdq      David Quigley      2212.06              1                   Michael Ambler
-----
91
```

This is a monthly report summarising Tinis HPC usage by research staff or students supervised by you. If any of these users you either do not recognise or they're no longer part of your group, please report this immediately via [bugzilla](#).

This is an automated email. Please do not reply.



Chiron

- ▶ Data analytics platform
 - Primarily funded for use by the CUSP and MathSys CDTs
 - Managed by the SC RTP
- ▶ Coprocessor system
 - 2 GPU nodes each with 2x NVIDIA Tesla K40
 - 2 Xeon Phi nodes each with 2 x Xeon Phi 7120P processors
- ▶ Fat node
 - 8-way fat node with 8TB RAM
- ▶ HADOOP cluster for offline analytics (560 cores)
- ▶ Streaming analytics cluster (640 cores)



Facilities

Managed Linux installation

Managed Linux

- ▶ Currently remote-manage 185 desktop PCs and 57 dedicated compute nodes over 7 departments
- ▶ Also accessible via remote-desktop (`godzilla`)
- ▶ Most machines contribute compute resource to the Cluster of Workstations (CoW)
(Largely superseded by Aprocrita?)
- ▶ 3800 unique users `active' in the last year
(Many of these make very little use of the system)



Managed Linux Software

- ▶ Based on SUSE Linux Enterprise Desktop 11 (SP4)
- ▶ Backported recent versions of key software
 - GNU compiler suite 4.9
 - Python 2.7 and 3.4
- ▶ Linux versions of site-licensed software
 - Matlab, Mathematica, Maple etc
- ▶ Licenses for additional commercial compiler suites
 - Intel compiler suite for C, C++, Fortran
 - Portland group compiler suite



Linux home directory

- ▶ File store (equivalent to ITS H: drive)
- ▶ Snapshots taken hourly, daily, monthly etc
- ▶ Disaster recovery backup to another building
- ▶ Contrary to popular belief, you **can** run software and packages from your home directory
(Lazily written installation guides often ignore this)




Group storage

- ▶ 1 PB of storage refreshed earlier this year
- ▶ Not backed up, but with some redundancy
- ▶ Accessible from all managed Linux machines and from the cluster log-in nodes, but not directly from the compute nodes.
- ▶ Used as working space for CoW jobs, large shared datasets etc
- ▶ Best regarded as a large file store than as a resource for data-intensive computing




Group storage

Group Storage Allocated (686 TB)



Group Storage Used (316TB)



Support

Getting help



- ▶ Online support request and tracking system
(c.f. ServiceNow from ITS)
- ▶ **Not** an automated system
 - There is a human system administrator reading the request – be polite!
 - Make an identifiable request or and/or ask a clear question
- ▶ **Not** a substitute for training in use of Linux
 - Many books and web-based guides available
 - May provide some training for new students in future
- ▶ **Not** there to help you use domain-specific software
 - Best we can do is put you in touch with other users





- ▶ Systems administrators are not psychic, even senior ones!

- ▶ Who are you?
 - Registering with just a first name is not helpful
 - Use your **university** email address to register
 - Introduce yourself – “I’m a new user from....”

- ▶ What are you trying to do and what have you tried so far?
 - Don’t assume we know what your research is about
 - We are not familiar with every piece of software you might use
 - What commands have resulted in errors, what is the error message?
 - Don’t assume the way you want something fixed is the best way





Bugzilla

- Tell us how important your request is so we can prioritise!
 - Is the issue blocking progress with your research?
 - Are you asking for something which is “nice to have” or essential?
 - How many people does the issue effect?
 - Only 3% of requests in the last year set anything but the default “P3” priority
 - Only 10% selected anything other than “normal” severity

Severity

	blocker	critical	major	normal	minor	enhancement	Total
P1	<u>1</u>	<u>5</u>	<u>7</u>	.	.	.	<u>13</u>
P2	.	.	<u>2</u>	<u>3</u>	.	.	<u>5</u>
P3	<u>4</u>	<u>6</u>	<u>24</u>	<u>688</u>	<u>2</u>	<u>15</u>	<u>739</u>
P4	.	.	<u>1</u>	<u>1</u>	.	.	<u>2</u>
P5	.	<u>1</u>	.	<u>1</u>	.	<u>2</u>	<u>4</u>
Total	<u>5</u>	<u>12</u>	<u>34</u>	<u>693</u>	<u>2</u>	<u>17</u>	<u>763</u>

Priority

[Bar](#) | [Line](#) | [Table](#) | [CSV](#)

Recent Developments

New regional (tier 2) computing facilities

New tier 2 systems

- ▶ In June 2016, EPSRC issued a call to refresh regional (tier 2) e-infrastructure for science broadly within their remit
- ▶ Diversity of hardware and workflow a key motivator
- ▶ 5 bids have been announced as successful
 - PETA-5 : £5M multi-architecture system
 - GW4 : £3M ARM cluster
 - Materials Modelling Hub : £4M x86 cluster
 - JADE : £3M *national* GPU cluster (Nvidia Pascal system)
 - HPC Midlands Plus : £3.2M x86 + OpenPower system
- ▶ The new Midlands system will be based at Loughborough, with Warwick committing 21% of the running and software support costs.




HPC Midlands Plus

- ▶ £3.2M is for hardware only
 - Will be primarily x86 with 11,000 to 13,000 cores depending on outcome of tendering exercise.
 - Will include fat GPU-equipped OpenPower nodes
- ▶ Access will require writing a minimal proposal
 - Light touch review process (volunteers needed)
 - Justify the need to run calculations at this scale
 - Details to be finalized
- ▶ Interested in “whacky runs” as well as large scale HPC jobs
 - Python managed HTC workflows banned from tier 1 HPC
 - Simulations dumping data to fat nodes for online processing



Recent Developments

Research software engineering (RSE)



Ian Holmes
@ianholmes

[Follow](#)

You can download our code from the URL supplied. Good luck downloading the only postdoc who can get it to run, though [#overlyhonestmethods](#)

RETWEETS
352

LIKES
134



8:52 am - 8 Jan 2013

[Reply](#) [Retweet](#) 352 [Like](#) 134 [More](#)

Research software engineers

- ▶ Significant initiative nationally to recognise importance of software.
- ▶ Many research papers represent what a single individual was able to get a piece of software to do, on a particular computer, once.
- ▶ Obvious consequences for reproducibility and duplication of effort.
- ▶ Science graduates on average have around 140 hours (i.e. 14 CATS) training in programming, insufficient to develop sustainable software to any meaningful standard.
- ▶ Need for RSEs to consult, train and support on development, deployment and documentation of research software.



RSE groups

- ▶ Warwick behind the curve
- ▶ Will invite RSEs to present to this forum over the next year
- ▶ Funding for 1 post to work 50% supporting HPC Midlands Plus, and 50% supporting a H2020 project in Maths
- ▶ Looking to grow this activity

Organisation	Group/Fellow
University of Bristol	Advanced Computing Research Centre
University of Cambridge	High Performance Computing Service
University College London	Research Software Development Group
Culham Centre for Fusion Energy	Data and Coding team (CODAS & IT)
University of Edinburgh	Edinburgh Parallel Computing Centre
Francis Crick Institute	Application Integration and Migration
Imperial College London	Research Software Engineering Community
ISIS	Mantid Group
University of Manchester	Research Software Engineering Group
University of Oxford	Advanced research Computing
University of Oxford	Research Software Developers Network
University of Sheffield	Research Software Engineering at Sheffield
University of Southampton	Research Software Group
STFC	Software Engineering Group

Recent Developments

Review of scientific computing at Warwick

Recommendations

- ▶ Disambiguate the SC RTP and the CSC
 - Recognise that facilities have outgrown the CSC
(Fraction of facilities use by CSC-affiliated academics is actually rather small)
- ▶ Improved communications
 - Better documentation and web presence (distinct from the CSC)
 - Clear definition of the responsibilities of the SC RTP vs ITS and departmental IT support
- ▶ Widen and support the user base (includes training)
 - Engage with data intensive research in the social sciences and beyond
- ▶ Clear guidance on how to cost use of facilities on research proposals
- ▶ Improved delivery of data storage
- ▶ Underwrite an RSE post within the SC RTP



Discussion?